

A BSSRDF Model for Efficient Rendering of Fur with Global Illumination

LING-QI YAN, University of California, Berkeley
WEILUN SUN, University of California, Berkeley
HENRIK WANN JENSEN, University of California, San Diego
RAVI RAMAMOORTHI, University of California, San Diego

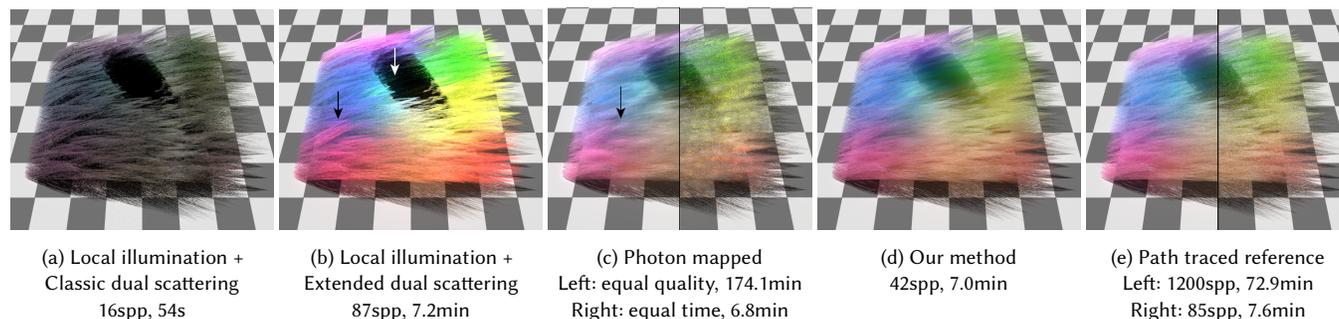


Fig. 1. Rendering of the *Pelt* scene using different methods, with a sphere casting a shadow onto it. The various colors are defined using a texture so that each fiber has a different color. (a) The classic dual scattering method fails to capture the scattered lobes from fur fibers, resulting in a dark appearance. (b) Our extended dual scattering handles scattered lobes, but is brighter and still produces hard and solid appearance, and does not have color bleeding effects. (c) Photon mapping is very inefficient for fur rendering, and is prone to overblur, as pointed out by the arrow. (d) Our method introduces a BSSRDF solution to the complex scattering, closely matching (e) the path traced reference, but is an order of magnitude faster.

Physically-based hair and fur rendering is crucial for visual realism. One of the key effects is global illumination, involving light bouncing between different fibers. This is very time-consuming to simulate with methods like path tracing. Efficient approximate global illumination techniques such as dual scattering are in widespread use, but are limited to human hair only, and cannot handle color bleeding, transparency and hair-object inter-reflection.

We present the first global illumination model, based on dipole diffusion for subsurface scattering, to approximate light bouncing between individual fur fibers. We model complex light and fur interactions as subsurface scattering, and use a simple neural network to convert from fur fibers' properties to scattering parameters. Our network is trained on only a single scene with different parameters, but applies to general scenes and produces visually accurate appearance, supporting color bleeding and further inter-reflections.

CCS Concepts: • **Computing methodologies** → **Reflectance modeling**;

Additional Key Words and Phrases: fur, global illumination, BSSRDF, neural network

ACM Reference format:

Ling-Qi Yan, Weilun Sun, Henrik Wann Jensen, and Ravi Ramamoorthi. 2017. A BSSRDF Model for Efficient Rendering of Fur with Global Illumination. *ACM Trans. Graph.* 36, 6, Article 208 (November 2017), 13 pages.
DOI: 10.1145/3130800.3130802

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2017 ACM. 0730-0301/2017/11-ART208 \$15.00
DOI: 10.1145/3130800.3130802

1 INTRODUCTION

Hair and fur are pervasive in recent movies and games. Significant recent progress has been made on physically-based local reflectance models for hair [Marschner et al. 2003] and fur [Yan et al. 2015, 2017] fibers. However, accurate rendering also requires global illumination for the diffusive and saturated appearance of the hair or fur volume. In fact, compared to local illumination that considers how light interacts inside individual fibers, global illumination caused by light scattering within the fur volume is usually brighter and more visible, composing the main part of the appearance.

However, global illumination rendering is slow. Therefore approximate methods are common, the most popular of which is the dual scattering technique [Zinke et al. 2008]. Dual scattering produces reasonably good results and allows real-time implementation.¹ It simplifies light scattering by assuming that all scattering events happen only along *main paths*, i.e. straight lines along incident directions. However, it is not applicable for fur rendering. This is mainly because the complex scattered lobes within individual fur fibers break the main path assumption. Even if we extend dual scattering to handle these scattered lobes, the result still doesn't match the path traced reference (Fig. 1). This will be analyzed in detail in Sec. 4. Furthermore, dual scattering does not support transparency of hair fibers as seen from the camera. Compared to path traced results, dual scattering still generates a hard and solid appearance (Fig. 1).

We develop a novel BSSRDF (Bidirectional Surface Scattering Reflectance Distribution Function) solution to global illumination, built

¹Currently, we focus on the offline rendering part of dual scattering as well as our method, though real-time implementation is a natural next step for our model.

upon the classic dual scattering, addressing many of dual scattering’s limitations. We analyze dual scattering in Sec. 4, and propose our model with three components: direct illumination from individual fur fibers, dual scattering to handle specular light transport, and BSSRDF for all other scattering events. In Sec. 5 and 6, we describe our BSSRDF model, and explain how to convert properties from fur fibers to BSSRDF parameters with the help of a multi-layer perceptron neural network (MLPNN). We validate our model in Sec. 7, and show close matches of our predicted renderings, compared with the path traced reference. Specifically, our model has these major advantages:

BSSRDF-based approximate global illumination. Our method is the first approximate global illumination model that is suitable for both hair and fur rendering.² We also provide the first empirical scheme of conversion from hair/fur parameters to BSSRDF parameters using a neural network. Our neural network is simple, consisting of only two hidden layers, and is fast to evaluate and easy to integrate into renderers for practical use. Furthermore, it only uses one scene for training with different parameters, and it generalizes well on others.

Color bleeding and accurate appearance. Dual scattering and other non-physically-based methods assume light transport only along main paths and assume local similarity with the fur fiber being shaded, resulting in opaque solid colors and BRDF style global illumination approximation. Our BSSRDF model is able to handle color bleeding from the fur volume for the first time, e.g. the color-filled shadow in Fig. 1 (c), and its general softer appearance. Moreover, in Sec. 7, we show that our model also generates much more accurate appearance, especially in terms of highlight and overall shading distributions, compared to the reference.

Further inter-reflections. Dual scattering works only with direct illumination and non-hair inter-reflections. The part of the indirect illumination from lit hair is missing with dual scattering. In contrast, our BSSRDF global illumination model (Sec. 4) naturally fits into offline renderers with further global illumination (Sec. 6), including approximate fur-to-fur inter-reflections for the first time.

Efficient performance. In addition to its accuracy, our method is also efficient. Thanks to the BSSRDF’s ability to “sum up” complex scattering events and higher order bounces, we’re usually able to achieve an order or magnitude speed up, compared with the path traced reference by Yan et al. [2017]. Detailed timing information is listed in Fig. 1 and Sec. 7.

2 RELATED WORK

In this section, we briefly describe related hair and fur local illumination models, and summarize existing global illumination methods for hair, both ad-hoc and physically-based. Note that, in some previous work, local illumination and global illumination for hair are called *single scattering* and *multiple scattering*. We avoid using these

terms throughout the paper, since multiple scattering events can happen even within single fur fibers.

Hair and fur reflectance models: Marschner et al. [2003] proposed the initial physically-based hair reflectance model. They approximate hair fibers as rough dielectric cylinders and use BCSDFs (Bidirectional Curve Scattering Distribution Functions) to describe light interactions within individual fibers. Based on whether light reflects (R) or transmits (T) when interacting with a cylinder, three types of lobes are formed: R , TT and TRT . Each lobe is separately evaluated longitudinally and azimuthally, as illustrated in Fig. 2. d’Eon et al. [2011] extended the Marschner model by fixing energy-conserving issues, such as adding a $TRRT$ lobe. Chiang et al. [2016] used one lobe to account for all high order internal reflections. Yan et al. [2015] proposed a physically accurate fur model, known as the double cylinder model, in which the cuticle, the cortex and the medulla are represented. The medulla scatters light, thus enriching the set of lobe types. Yan et al. [2017] further simplifies these types, so that only two scattered lobes TT^s and TRT^s are added to the previous R , TT and TRT lobes, to model the complex scattering behavior from the medulla, as labeled in Fig. 2 (c).

Hair global illumination methods: Accurate global illumination requires simulating actual light bouncing between hair fibers. Moon et al. [2006] extended photon mapping to store and query photons within the hair volume as light bounces inside. Hery et al. [2012] and d’Eon et al. [2013] proposed different importance sampling schemes for hair BCSDFs to accelerate the convergence of path traced global illumination.

Most non physically-based approximate global illumination methods treat hair fibers as semi-transparent. To compensate for light transmittance through fibers to the shading point, shadow map based methods [Lokovic and Veach 2000; Sintorn and Assarsson 2009; Yuksel and Keyser 2008] accumulated transparency of hair fibers to approximate the optical thickness of hair fibers that the light goes through. To enable transparency looking from the camera, alpha blending based methods [Endernton et al. 2011; Sintorn and Assarsson 2008, 2009; Yu et al. 2012] proposed different ways to approximate the back-to-front blended self-occlusion effects. While generating plausible results, none of these methods are physically-based, and they have not been shown to be applicable for accurate fur reflectance models.

The only physically-based approximation to hair global illumination is the dual scattering approximation method [Zinke et al. 2008]. It assumes that scattering events happen along the main path – the light reaches the shading point by penetrating through the hair volume in a straight line (global scattering), and scatters along the camera path into the hair volume, then back to the shading point (local scattering). Though successfully used in practice, dual scattering cannot be applied for fur rendering (Sec. 4). Furthermore, it doesn’t account for transparency or color bleeding.

Subsurface scattering: The idea of subsurface scattering was introduced to computer graphics by Blinn et al. [1982]. Different approaches were applied to solve it, such as path tracing [Jensen et al. 1999], photon mapping [Dorsey et al. 1999] and scattering equations [Pharr and Hanrahan 2000], but were too costly to make it practical. Jensen et al. [2001] introduced a dipole solution to solve subsurface scattering within translucent materials efficiently, which was later

²In Yan et al. [2017], the concepts of hair and fur are unified. Complex scattering structures have been discovered in human hair fibers as well, and the benefit of these structures has also been demonstrated in hair rendering.

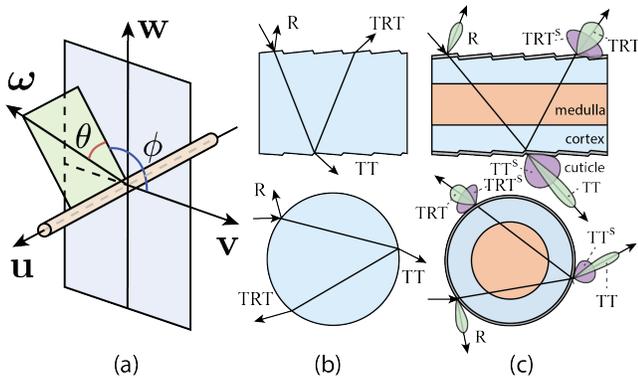


Fig. 2. (a) Longitudinal-azimuthal parameterization for hair/fur fibers. Any direction ω can be parameterized into θ in the plane spanned by ω and the cylinder axis u , and ϕ orthogonal to the plane. (b) Longitudinal and azimuthal lobes of the Marschner hair model. (c) Lobes of the fur model by Yan et al. [2017].

extended to extended dipole [Donner and Jensen 2008], quantized diffusion [D'Eon and Irving 2011], anisotropic dipole [Jakob et al. 2010] and directional dipole [Frisvad et al. 2014] for better accuracy. These advanced methods could in theory improve accuracy, but in practice, we find that using the original dipole solution by Jensen et al. [2001] already gives good results.

Participating media parameter conversion: To represent complex appearance, previous work explored ways to convert from micro structures to scattering parameters. Granular materials rendering [Meng et al. 2015; Müller et al. 2016] uses BSSRDF models for fast approximation. However, the derivation assumes isotropic distributed grains, but hair and fur are usually highly disciplined and anisotropic; thus it is not directly applicable. In cloth rendering, it is a common approach to abstract complex fibers into general participating media [Khungurn et al. 2015; Zhao et al. 2011, 2016]. They depend on simple reflection models of cloth fibers, and render them as high resolution volume grids. To convert from fibers' properties to parameters of participating media, the most popular model is the microflake model [Heitz et al. 2015; Jakob et al. 2010], where the volume is assumed to be filled with randomly oriented flakes according to some distribution. However, we will show that the microflake model does not produce good results in our case in Sec. 4, and so we use neural networks.

3 BACKGROUND

In this section, we briefly recap some background knowledge. We first describe hair and fur BCSDf models for local illumination from individual hair/fur fibers. Then we talk about the dual scattering approximation for hair global illumination. Finally, we introduce the general subsurface scattering model and its dipole solution. We propose our model in Sec. 4.

3.1 Hair and fur BCSDf models

Hair and fur reflectance models use dielectric cylinders to represent hair and fur fibers. Light can get reflected (R) or transmitted (T) each time it interacts with the surface of a cylinder, representing the hair or fur *cuticle*. Inside the cylinder known as *cortex*, it gets

Table 1. Parameters used for hair model (first 4 parameters) and fur model in Yan et al. [2017] (all 9 parameters).

Parameter	Definition
η	refractive index of cortex
α	scale tilt for cuticle
β_m	longitudinal roughness of cuticle (stdev.)
$\sigma_{c,a}$	absorption coefficient in cortex
β_n	azimuthal roughness of cuticle (stdev.)
κ	medullary index (rel. radius length)
$\sigma_{m,a}$	absorption coefficient in medulla
$\sigma_{m,s}$	scattering coefficient in medulla
g	anisotropy factor of scattering in medulla
l	layers of cuticle

partially absorbed thus producing color. The differences are that, hair models use one cylinder (Fig. 2 (b)), producing unscattered lobes $p \in R, TT, TRT$. In contrast, the fur model [Yan et al. 2017] (Fig. 2 (c)) has an additional inner cylinder or *medulla* that scatters light, adding scattered lobes TT^s and TRT^s to p . However, more parameters are required to describe the model. We list the parameters used in both Marschner et al. [2003] and Yan et al. [2017] in Table 1.

The interactions with light between hair and fur fibers are described using BCSDFs, often represented by the longitudinal-azimuthal (θ, ϕ) parameterization as illustrated in Fig. 2,

$$L_r(\theta_r, \phi_r) = \int_{-\pi}^{\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} L_i(\theta_i, \phi_i) S_c(\theta_i, \theta_r, \phi_i, \phi_r) \cos^2 \theta_i d\theta_i d\phi_i, \quad (1)$$

where S_c is the BCSDf, (θ_i, ϕ_i) and (θ_r, ϕ_r) are the incident and outgoing directions, L_i and L_r are the incoming and outgoing radiance.

Similar to different lobes in BRDFs for surface reflectance, different hair and fur models represent their BCSDFs using a series of lobes representing different types of light interactions. These lobes are often factored into a product of M and N lobes, describing the light's interactions on the longitudinal and azimuthal sections separately. These factored lobe-based BCSDFs share the same representation as

$$\begin{aligned} S_c(\theta_i, \theta_r, \phi_i, \phi_r) &= \sum_p S_p(\theta_i, \theta_r, \phi_i, \phi_r) / \cos^2 \theta_d \\ &= \sum_p M_p(\theta_h) \cdot N_p(\phi; \eta') / \cos^2 \theta_d, \end{aligned} \quad (2)$$

where $\theta_h = (\theta_r + \theta_i)/2$ and $\theta_d = (\theta_r - \theta_i)/2$ are the longitudinal half angle and difference angle, and $\phi = \phi_r - \phi_i$ is the relative outgoing azimuth in $[-\pi, \pi]$. The azimuthal section is always treated as circular, while the effect of an inclined longitudinal incident angle is incorporated with a virtual IOR η' .

Specifically, Yan et al. [2017] proposes a simple fur reflectance model, where the TT and TRT paths generate additional scattered lobes TT^s and TRT^s when they pass the medulla, or the inner cylinder. So, there are 5 lobes in their model (Fig. 2 (c)). The scattered lobes are also longitudinal-azimuthal separated, and both of M and N are calculated by querying precomputed scattering profiles.

Our method builds upon the fur model of Yan et al. [2017] with unscattered lobes and scattered lobes. Yan et al. [2017] indicate that a small medulla can also improve the appearance of human

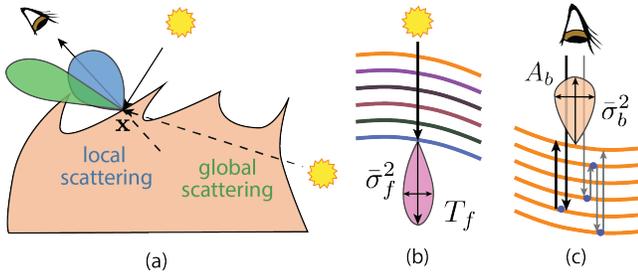


Fig. 3. (a) Illustration of dual scattering for approximate global illumination. (b) Attenuation and spread computation for global scattering. (c) Attenuation and spread computation for local scattering.

hair rendering, and unify hair and fur reflectance. Therefore, our global illumination method also benefits hair rendering, and can be considered a general approach for hair and fur.

3.2 Dual scattering approximation

Dual scattering approximates global illumination effects within the hair volume at a shading point x as a combination of two components: global scattering and local scattering, illustrated in Fig. 3 (a). The global scattering approximates how much light arrives at x after penetrating through n fibers along the light path. If the light comes from behind x , it is directly seen by the camera, forming a *global scattered lobe*. If not, the arrived light will be added to direct illumination, going into the hair volume, scattering inside, then going back, forming a *local scattered lobe*.

Dual scattering begins with simplifying the complicated longitudinal lobes for each hair fiber as one forward lobe and one backward lobe, as illustrated in Fig. 4 (a). It first pre-computes the averaged forward/backward scattering intensity $\bar{a}_{f|b}$ as

$$\bar{a}_{f|b}(\theta_i) = \frac{1}{\pi} \iint_{\Omega_{f|b}} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} S(\theta_i, \phi_i, \omega_r) \cos(\theta_r) d\phi_i d\omega_r, \quad (3)$$

where $\omega_r = (\theta_r, \phi_r)$ is the outgoing direction to the forward or backward hemisphere $\Omega_{f|b}$.

Apart from the averaged forward/backward intensities, the forward/backward lobes also require their averaged longitudinal variances $\bar{\beta}_{f|b}^2$. Since the TT lobe is mostly forward and the TRT lobe is mostly backward, the variances directly take their squared roughness $\bar{\beta}_f^2 = \beta_{TT}^2$ and $\bar{\beta}_b^2 = \beta_{TRT}^2$. The R lobe is simply ignored.

To evaluate global and local scattered lobes, the key is to compute how the light attenuates and spreads longitudinally along a main path. Figures 3 (b) and (c) illustrate the computation. Dual scattering assumes that the azimuthal scattering is complicated, and it always becomes isotropic.

For the global scattered lobe, dual scattering calculates the attenuation of the light T_f reaching x after forward scattering through n fibers by taking the sequential product of their averaged forward attenuations. And the spread variance $\bar{\sigma}_f^2$ reaching x is computed by accumulating the pre-computed variances of all n fibers. Figure 3 (b) illustrates this idea.

The local scattered lobe is approximated similarly, as shown in Fig. 3 (c). The camera path goes into the hair volume, going forward through i hair fibers, scattering back once, then going forward through the previous i fibers until reaching the shading point x

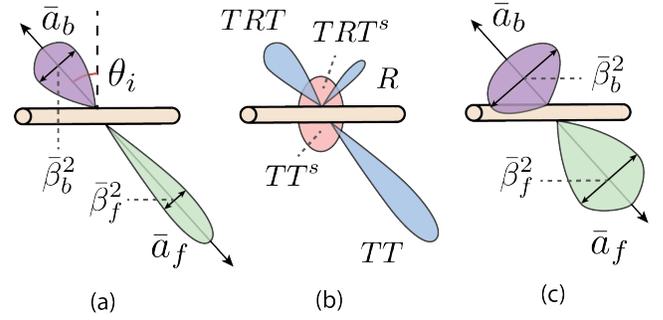


Fig. 4. (a) Averaged forward/backward scattering intensities $\bar{a}_{f|b}$ and variances $\bar{\beta}_{f|b}^2$ for a single hair fiber. (b) All 5 longitudinal lobes in Yan et al. [2017]. Note the scattered lobes TT^s and TRT^s that are too smooth and wide to follow the main path assumption. (c) Simple extension of dual scattering by calculating the contribution of scattered lobes to the forward/backward scattered lobes anyway.

again. By summing up all possible paths with varying $i \geq 1$, the local attenuation with one backward scattering event A_1 can be computed. In order that the camera path finally returns to x , backward scattering can also happen any odd number of times i.e. 3, 5, ... Dual scattering also computes A_3 , and ignores higher-ordered scattering, so that the total attenuation for the backward scattered lobe is $A_b = A_1 + A_3$. The spread variance $\bar{\sigma}_b^2$ is the averaged accumulated variance along each possible path, weighted by its attenuation.

With global and local scattered lobes computed, the final approximate global illumination is the sum of both lobes, added to the hair BCSDf model. Specifically, the local scattering forms a lobe:

$$S_b(\theta_i, \theta_r, \phi) = \frac{I_b(\phi)}{\pi \cos^2 \theta_i} \cdot d_b A_b \cdot G(\theta_r + \theta_i; \bar{\sigma}_b^2), \quad (4)$$

where d_b is the backward scattering density constant, usually set between 0.6 and 0.8, while $G(\mu; \sigma^2)$ is a Gaussian with mean μ and variance σ^2 . $I_b(\phi)$ is a binary backward hemisphere indicator which is 1 when $\phi \in [-\pi/2, \pi/2]$ and 0 elsewhere.

The global scattering forms another lobe:

$$S_f(\theta_i, \theta_r, \phi) = \frac{I_f(\phi)}{\cos^2 \theta_i} \cdot d_f A_f \cdot \sum_p \left(G(\theta_r + \theta_i; \bar{\beta}_p^2 + \bar{\sigma}_f^2) N_p^f \right), \quad (5)$$

where d_f is the forward scattering density constant, usually set between 0.6 and 0.8. $I_f(\phi)$ is a binary forward hemisphere indicator which is 0 when $\phi \in [-\pi/2, \pi/2]$ and 1 elsewhere. $N_p^f = \frac{1}{\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} N_p(\phi - \phi'; \eta') d\phi'$ is the averaged azimuthal lobe p within the front hemisphere.

The dual scattering has made many assumptions, the most important of which is the main path assumption. In Sec. 4, we demonstrate that this assumption is only reasonable for unscattered lobes. For scattered lobes, it will fail. So, in our model, we only use dual scattering as a component to handle paths consisting of unscattered lobes (TT and TRT , as in hair rendering).

The failure of dual scattering is mainly because the scattered lobe is too wide and uniform to be considered along the main path, as illustrated in Fig. 4 (b). We will discuss a dual scattering extension in Sec. 4.2 that results in wide forward/backward lobes (Fig. 4 (c)).

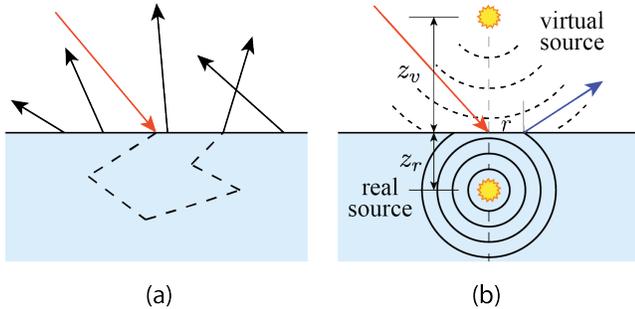


Fig. 5. (a) Illustration of BSSRDF, where the light exits at different points than the incident point. (b) Dipole method approximating BSSRDF by summing up contribution from a real source and a virtual source to an exiting point at distance r .

3.3 Subsurface scattering

Instead of being directly reflected at the shading point, light may go into the subsurface, scatter, and exit at a different point, as shown in Fig. 5 (a). Note that, the idea of subsurface scattering only applies to regular surfaces, and has not previously been attempted for hair and fur rendering. One of our key innovations is to use subsurface scattering ideas for global illumination in hair and fur, as will be described in detail in Sec. 4. In this subsection, we introduce subsurface scattering on regular surfaces.

The subsurface light transport is often represented using BSSRDFs, leading to a generalized rendering equation

$$L_o(x_o, \omega_o) = \int_{\mathcal{A}} \int_{\Omega} L_i(x_i, \omega_i) S_{ss}(x_i, \omega_i; x_o, \omega_o) (\mathbf{n} \cdot \omega_i) d\omega_i d\mathcal{A}(x_i), \quad (6)$$

where S_{ss} is the BSSRDF, extending a BRDF with different incident and outgoing positions. \mathcal{A} is the area associated with the incident position x_i .

To represent BSSRDFs, material properties that are responsible for the scattering behavior must be defined: σ_a is the absorption coefficient, σ_s is the scattering coefficient, $\sigma_t = \sigma_a + \sigma_s$ is the extinction coefficient, and $\alpha = \sigma_s / \sigma_t$ is the albedo. For anisotropic scattering, an anisotropy factor $g \in [-1, 1]$ is defined, resulting in reduced scattering coefficient $\sigma'_s = (1 - g)\sigma_s$, reduced extinction coefficient $\sigma'_t = \sigma_a + \sigma'_s$ and reduced albedo $\alpha' = \sigma'_s / \sigma'_t$.

While the BSSRDF is often difficult to calculate, Jensen et al. [2001] proposed a dipole method to solve the multiple scattering part of it. As illustrated in Fig. 5 (b), the dipole method assumes local flatness, putting a real point source beneath the incident x_i and a virtual point source above it. Then the diffuse reflectance at the outgoing position x_o of distance r is the contribution of these dipole sources:

$$R_d(r) = \frac{\alpha' z_r (1 + \sigma_{tr} d_r) e^{-\sigma_{tr} d_r}}{4\pi d_r^3} - \frac{\alpha' z_v (1 + \sigma_{tr} d_v) e^{-\sigma_{tr} d_v}}{4\pi d_v^3}, \quad (7)$$

where $z_r = 1/\sigma'_t$ and $z_v = -z_r(1 + 4A/3)$ are the positive and negative z -coordinates of the real and virtual point sources, respectively. Here $A = (1 + F_{dr})/(1 - F_{dr})$, and F_{dr} is a refractive index related variable. $d_r = \sqrt{r^2 + z_r^2}$ and $d_v = \sqrt{r^2 + z_v^2}$ are distances from x_o to the sources. $\sigma_{tr} = \sqrt{3\sigma_a\sigma'_t}$ is the effective extinction coefficient.

The BSSRDF due to multiple scattering, a.k.a. the diffusion term is then

$$S_d(x_i, \omega_i; x_o, \omega_o) = \frac{1}{\pi} F_t(\eta, \omega_i) R_d(\|x_i - x_o\|) F_t(\eta, \omega_o), \quad (8)$$

where F_t terms are for Fresnel transmission.

Apart from the multiple scattered diffusion, to complete the BSSRDF, a single scattering term $S^{(1)}(x_i, \omega_i; x_o, \omega_o)$ is added to account for cases where only one scattering event happens. Note that the single-multiple scattering separation is approximate. And in our model, we use a different separation scheme, as will be introduced in Sec. 4.

To accelerate rendering using the dipole model, Jensen et al. [2002] later proposed a two-pass algorithm. In the first pass, they uniformly distribute sample points across the surface of the translucent object. Each sample point's irradiance is evaluated through a path tracing process. The second pass is a traditional path tracing process, where nearby sample points are queried for their contribution to each shading point. To accelerate this process for locating points, a hierarchical octree structure is built on top of these sample points, where each node represents a single sample point, with its irradiance value and position the average of its child nodes. The octree structure is traversed top-down to perform quick rejection of samples that are far enough from the shading point.

4 MOTIVATION AND OVERVIEW

In this section, we briefly analyze the limitations of dual scattering to motivate our method. Then we provide a high level overview of our full global illumination model, specifying different components it consists of.

4.1 Limitations of dual scattering

In dual scattering, there is an important assumption that the longitudinal lobes from hair BCSDFs (thus the averaged forward and backward lobes) are sharp. This is why it is reasonable that the light only transports along main paths. While generally true for human hair, when applied to animal fur, the model will break. The medullas within fur fibers produce smooth and diffusive scattered lobes, which do not exist in previous hair models and are not suitable to be abstracted using a forward and a backward lobe, as shown in Fig. 4.

Even if we do assume light scattering only along main paths, the rendering results in Figs. 1 and 17 indicate that dual scattering still cannot get us the correct light spread or color bleeding around the shading point \mathbf{x} . This is because dual scattering always adds approximate global illumination only to \mathbf{x} . Intuitively, the difference is analogous to BRDFs vs. BSSRDFs. This observation motivates us to model global illumination using BSSRDFs, as will be elaborated in Sec. 4.3 and 5.

Another limitation of dual scattering is that it is unclear how to handle multiple bounces. One possible approach is to trace secondary rays through the hair volume until they hit non-hair geometry. All hair intersections along the secondary rays are considered as a part of the global scattering component of dual scattering. However, this approach usually produces identical images, compared to the classic dual scattering method without multiple bounces. This is because the secondary ray's contribution will quickly fall off due to the hair fibers it penetrates, and the absorbed incident radiance before reaching the non-hair geometry. Because of this, we typically

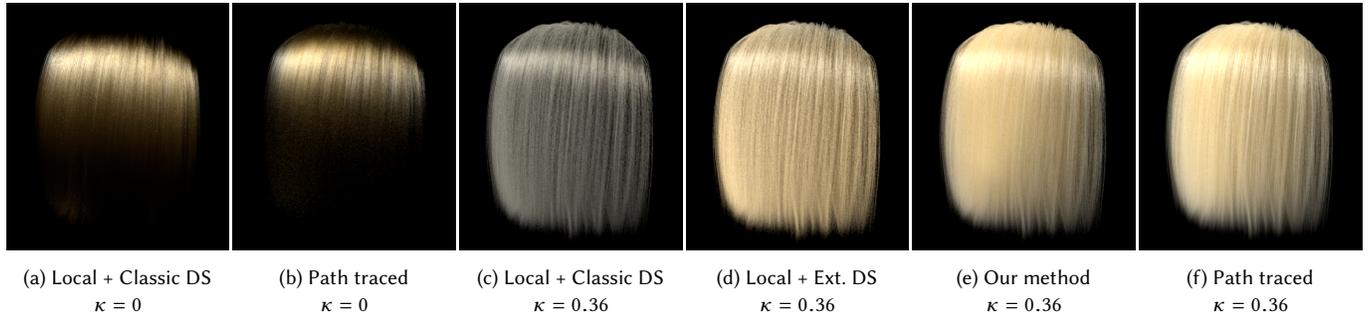


Fig. 6. Classic dual scattering (a) works well on traditional hair models without medulla compared with path traced reference (b), but cannot capture scattered lobes from fur models with medulla (c). Our extended dual scattering (d) approximately handles the medulla and resolves the energy loss, but still doesn't match path traced reference (f) as well as our method (e).

use standard extended dual scattering for efficiency in the rest of the paper for comparison, unless multiple bounces are clearly stated, as in Figs. 16 (c) and 18 (a).

4.2 Extending dual scattering

Since dual scattering only works for specular unscattered lobes, simply applying it to these lobes will result in energy loss from scattered lobes (see Figs. 1 (a) and 6 (c)). To guarantee fair comparison with our approach later, we extend dual scattering to approximately handle the scattered lobes.

Our extension is to simply average the attenuation and spread of all lobes, regardless of whether they are spread or not (Fig. 4 (c)). We calculate the averaged forward/backward attenuation $\bar{a}_{f|b}$ using Eqn. 3, but using all lobes including TT^s and TRT^s . We average the lobe-weighted average spread similar to Sadeghi et al. [2010] as:

$$\bar{\beta}_{f|b} = \frac{\int_{\Omega_{f|b}} \sum_{p \in \{R, TT, TRT, TT^s, TRT^s\}} S_p \beta_p d\omega}{\int_{\Omega_{f|b}} \sum_{p \in \{R, TT, TRT, TT^s, TRT^s\}} S_p d\omega}, \quad (9)$$

assuming the standard variances of the TT^s and TRT^s lobes as $\beta_{TT^s} = \beta_{TRT^s} = \pi/4$, since they are approximately uniformly distributed in the longitudinal section. The spread is pre-computed by uniformly sampling θ and ϕ and numerically calculating the integrals.

Note that, our extension to dual scattering is a bold assumption. It will result in unreasonably wide forward and backward lobes, but at least allows for energy conservation (See Fig. 4). Moreover, it still makes the main path assumption, thus resulting in no color bleeding effects, as pointed out in Fig. 1 (b). Furthermore, the brightness of the extended dual scattering is not stable (with fixed $d_f = d_b = 0.7$). Sometimes it will generate much brighter appearance than the reference (Fig. 1) while sometimes being darker (Fig. 16).

We compare the classic and our extended dual scattering in Fig. 6 (a)-(d). We first use the classic dual scattering on human hair without medulla (a) together with local illumination. Since no scattered lobes exist, this is a case where the classic dual scattering works well, compared to the reference (b). Then we apply the classic dual scattering and our extended dual scattering with a medulla of width $\kappa = 0.36$ (c)-(d), also with local illumination, and compare them with our method (e) and reference (f). We find that the classic dual scattering suffers from severe energy loss. The extended dual scattering alleviates this issue, but generates a flat appearance. In contrast, our

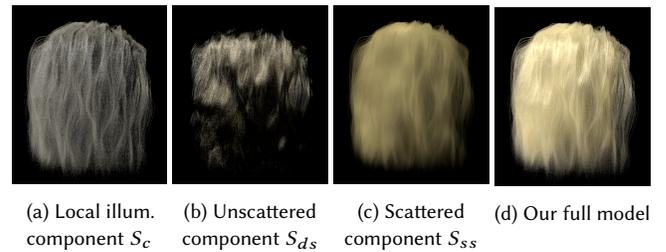


Fig. 7. Decomposed components in our model. (a) Local illumination component S_c , including all 5 lobes: R , TT , TRT , TT^s and TRT^s . (b) Unscattered component S_{ds} captured by classic dual scattering, including R , TT and TRT lobes. (c) Scattered component S_{ss} represented using BSSRDF, including all light interactions that are related to TT^s and TRT^s lobes. (d) Our full model with all components.

BSSRDF model, introduced next, generates a close match with the reference, as will be introduced next.

4.3 Model overview

To break the main path constraint, we represent our global illumination model as a combination of a BCSDf and a BSSRDF. The BCSDf also contains two components: direct illumination and dual scattering for multiple scattering from only specular lobes. The BSSRDF is responsible for all other multiple scattering events.

The direct illumination component describes light interaction with single hair or fur fibers. It is thus the hair or fur's BCSDf S_c in Eqn. 2. We name this part as the *local illumination component*.

Another part is the dual scattering of specular lobes S_{ds} . This is part of the global illumination. Since there are no scattered lobes, the main path assumption still holds. Specifically, we convert only R , TT and TRT lobes into forward and backward lobes using Eqn. 3, and use them to compute dual scattering's contribution. We define this part as the *unscattered component*. This part is slightly more accurate than the classic dual scattering method, since we also take the R lobe into account.

For the rest of the energy, at least one scattered lobe contributes. Since the scattered lobes are usually smooth and more isotropic than unscattered specular lobes, significant scattering happens. Thus, we use subsurface scattering to capture the scattering effects. Specifically, given the parameters of hair or fur, we convert them to subsurface parameters, then use the dipole method for rendering. We

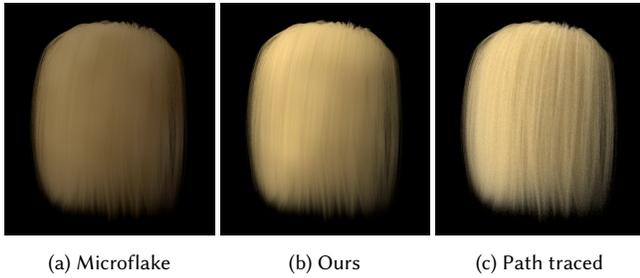


Fig. 8. Rendering of scattered component using different methods. Our empirical BSSRDF model fits the path traced reference much closer, compared to the physically derived microflake model.

name this part as the *scattered component*. Note that, for hair which doesn't have a medulla, this component becomes zero. In this case, our method reduces to classic dual scattering.

Our final model can be represented as:

$$S_{\text{BCSDF}} = S_c(\omega_i, \omega_r) + S_{ds}(\omega_i, \omega_r), \quad (10)$$

$$S_{\text{BSSRDF}} = S_{ss}(x_i, \omega_i, x_o, \omega_r), \quad (11)$$

where S_c is individual hair or fur fiber's BCSDF. S_{ds} is the dual scattering approximated unscattered component, which is essentially another BCSDF. S_{ss} is our estimated BSSRDF scattered component.³ For an intuitive illustration of different components, we visualize the decomposed appearance of each component in Fig. 7.

Though the idea of using a BSSRDF model to represent the scattered component may sound straightforward, it is far from easy, both in theory and in practice. First, BSSRDFs work only with actual surfaces. However, when it comes to hair or fur fibers, there is no concept of a surface. For this reason, distributing dipole samples for hair and fur and calculating associated areas with them is difficult. Second, there is no existing theory to convert physically based hair and fur properties into parameters of BSSRDFs. Third, the dipole solution to BSSRDFs only transports light locally in the shading point's neighbourhood. So, indirect lighting from the same hair or fur volume cannot be accounted for, resulting in dark areas with energy loss. In Sec. 5 and 6, we elaborate how to deal with all these difficulties.

4.4 Physically-based derivation of BSSRDF parameters

Theoretically, there is no existing BSSRDF model that is suitable to represent hair and fur volumes, since the hair and fur fibers define rather anisotropic scattering behavior. As introduced in Sec. 2, the closest match is the microflake model. Deriving a physically based model based on the microflake theory for fur fibers is not impossible. However, we find it does not work well in practice, and we turn to a data-driven approach to solve the problem, as will be introduced in the next section. Here we assume that the microflake model is the natural first step to try, so we describe it below and compare it with our BSSRDF model in Fig. 8.

We think of each hair/fur fiber as a series of small "flakes" with their BCSDFs as phase functions [Jakob et al. 2010]. However, this simple conversion is not well-defined in practice, mainly because of the high complexity of hair/fur fibers. The BCSDF of a hair/fur

³The BCSDF and BSSRDF components are defined in different domains, and they are integrated using Eqns. 1 and 6, respectively.

fiber consists of 5 lobes, R , TT , TRT , TT^s and TRT^s , which all have rather complicated shape and large variation with the incoming light direction. However, in microflake theory, each microflake is completely opaque and mirror reflective. Nevertheless, it is still possible for us to derive microflake parameters by assuming local similarity, as discussed in the supplementary material.

We apply the derived parameters to render an actual scene, as shown in Fig. 8. As expected, the result is still far from the path traced reference. The different shading distribution demonstrates the inability of individual flakes to represent hair/fur fiber segments, as analyzed above. Moreover, we also find a color difference as compared to the reference. This is because the color mainly comes from the absorption coefficient σ_a , which is usually two orders of magnitude smaller than the scattering coefficient σ_s . So, a slight inaccuracy in σ_a will result in significant color difference.

5 BSSRDF APPROXIMATION FOR SCATTERED COMPONENTS

As discussed in Sec. 4.4, since deriving a practical physically based model to estimate the dipole parameters from the hair/fur parameters is challenging, we turn to a data driven approach. We propose a neural network structure to solve the parameter conversion problem. We chose a neural network over other methods because the dimensionality of the input and output spaces of the problem is high and the conversion function is likely to be non-linear, making simple fitting methods impractical.

In this section, we first generalize the dipole model for hair/fur geometry in Section 5.1 to be able to use the model in the first place. Then, in Section 5.2, we describe our neural network architecture for the parameter conversion problem. In Section 5.3, we describe how to train the neural network efficiently.

5.1 Generalizing BSSRDF for Fur/Hair Geometry

We extend the 2-pass algorithm introduced in Section 3.3 to enable hair/fur models in dipole rendering. The only missing component here is how to perform point sampling on a hair/fur geometry.

Figure 9(a) shows a hair/fur geometry representation, where each hair/fur fiber is represented by an array of vertices v_0, \dots, v_3 with associated radiuses r_0, \dots, r_3 . The geometry represented by each 2 consecutive vertices is a frustum. To generate a sample point, we first randomly select a frustum from all frustums in the hair/fur geometry with a probability proportional to its side surface area. Then, a sample point can be obtained by a uniform sampling on the selected frustum's side surface as in Fig. 9(b). Since the sample points generated this way may be highly occluded from the environment as illustrated in Fig. 9(c), we remove those points that have zero estimated irradiance in the first pass in the octree construction process for efficiency.

5.2 Neural Network for Parameter Conversion

In this section, we now describe our neural network architecture to estimate the parameters of a dipole model given the parameters of a fur/hair model. The appearance of a fur/hair model is decided by the 9 material parameters listed in Table 1. From our experiments, we find that only 5 parameters, κ , $\sigma_{c,a}$, $\sigma_{m,s}$, η and l out of all 9 parameters have observable impact on the appearance of the scattered component. Intuitively, this is because g and α are often

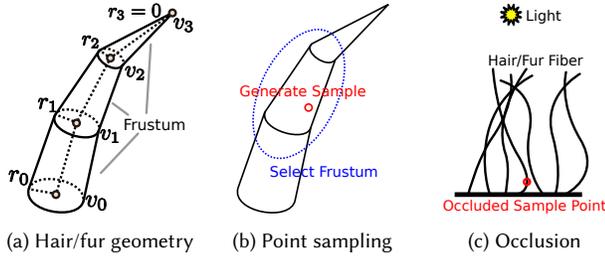


Fig. 9. Sample placement for hair/fur geometry.

very close to zero in realistic fur/hair models, so they only have very subtle impact on the scattered component. The roughness parameter of unscattered lobes β is small compared with the roughness of scattered lobes, so its effect is minimal in the scattered component where each path has at least one scattered interaction. We don't use $\sigma_{m,a}$ as input because it is correlated with κ and $\sigma_{c,a}$.

On the other hand, the parameters of a dipole model include σ_a and σ_s as we introduced in Section 3.3. To enhance the flexibility of the dipole model, we add another brightness scaling parameter w , which is multiplied with the rendered dipole component at the end. Note that for the dipole model, the anisotropy parameter g and σ_s are correlated. Therefore, we only change σ_s as in the dipole model and always set g to zero. We do not have any special restriction on the value of w in order to have more flexibility in the model.

Given the high dimensionality of the parameter spaces, we use a multi-layer perceptron neural network (MLPNN) as our model for the parameter conversion problem. Figure 10 shows our model structure. For the input parameters, we first apply a preprocessing step to map them into an appropriate range that is easy for neural networks to train. The preprocessing function for each input is listed in the figure. Then, we feed the 5 preprocessed input parameters, $X = \{x_0, \dots, x_4\}$ into a MLPNN with 2 10-node fully connected hidden layers using tanh activation function to produce 3 output parameters, $Y = \{y_0, y_1, y_2\}$. Finally, we apply a post-process to convert the output to dipole model parameters σ_a , σ_s and w as follows:

$$\sigma_a = 2^3 \tanh(y_0) + 5, \quad \sigma_s = 2^3 \tanh(y_1) + 5, \quad w = 5 \tanh(y_2) + 5. \quad (12)$$

Note that the mapping we choose in post-processing is to fit our training scene settings, which we will elaborate in Section 5.3.

We chose a very simple neural network structure instead of deep neural networks mainly for two reasons. First, the parameter conversion function should be very smooth given that the appearances of both models vary smoothly with changing input parameters. Therefore, using a simple structure is sufficient and avoids over fitting. Second, the neural network needs to be evaluated at each shading point during rendering to support heterogeneous fur models. Using a simple structure makes sure that the evaluation is relatively cheap compared to ray tracing. We validate in Section 6 that the neural network evaluation time is less than 10% of the total render time. Our neural network also has negligible memory footprint, since it only has about 200 parameters.

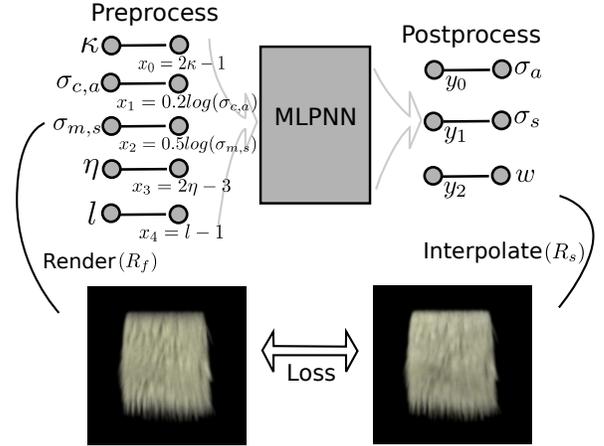


Fig. 10. Our neural network structure for the parameter conversion problem.

5.3 Training the Neural Network

To train the neural network, we need to quantitatively measure the quality of each input-output pair it generates. To do this, as illustrated in Fig. 10, we compare the rendered image using the hair/fur model with input parameters and the rendered image using the dipole model with output parameters under the same scene settings. The resemblance between the 2 images should be a good indicator of how good the parameter conversion predicted by the neural network is.

Scene Settings. We use a single scene to train our neural network. Figure 10 shows an example of renderings produced by an input-output pair using the training scene. The model in our training scene is a piece of fur pelt directly facing the camera. We use the same camera and light settings and only change the material properties of the model throughout the training process. We show in Section 7 that although our model is only trained on one scene, it generalizes to many other scenes with different geometry, lighting, etc. This is because the parameter conversion function is not sensitive to the scene settings, and is an intrinsic property of the materials from the 2 models (hair/fur and dipole) that we try to match. Here, we only use the training scene to find a good approximation to the parameter conversion function.

The output dipole parameters of the neural network fit the training scene well. However, if we scale the training scene by a factor, σ_a and σ_s need to be scaled accordingly to obtain correct results. In Section 6, we describe how to do this in scenes with different scales.

Loss Function. Suppose X is an input parameter set and $R_f(X)$ is its rendered image using the hair/fur model with only the scattered component. Recall that the unscattered component is handled separately using the traditional dual scattering in our model as introduced in Section 4.3. Therefore, we use path traced scattered components as the reference image for the dipole model to fit. $Y = f(X)$ is the output parameter set predicted by the network, where f is the parameter conversion function approximated by the neural network. $R_s(Y)$ is the rendered image using the dipole model and the output parameter set. Our loss function is defined as:

$$\text{Loss} = \gamma L_1(R_f(X), R_s(Y)) + L_s(R_f(X), R_s(Y)), \quad (13)$$

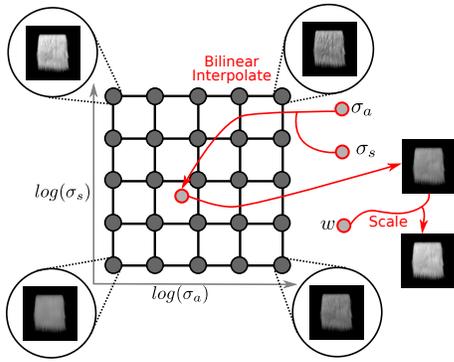


Fig. 11. We approximate R_s using bilinear interpolation. Before training the neural network, we store a 2D grid of pre-rendered images with varying σ_a and σ_s in log space. Given the 3 dipole parameters, σ_a , σ_s and w , we first use σ_a and σ_s to perform bilinear interpolation to obtain the appearance of the rendered image. Then, we scale the interpolated image by w to obtain the final approximated rendering.

where L_1 is the L_1 norm between 2 images. $L_s(R_f(Y), R_s(X)) = 1 - \text{SSIM}(R_f(Y), R_s(X))$ is the structural loss term, where SSIM is the structural similarity index [Wang et al. 2004] between 2 images. γ is a parameter to weight the relative impact. In our training, we set $\gamma = 0.1$ to stress more on the structural loss term. This is because the structural loss term has much more impact on the overall visual quality than the L_1 term.

Approximate rendering using bilinear interpolation. For efficient training, R_s needs to have a fast evaluation method and be differentiable in order to do gradient back-propagation. Finding the gradients can be difficult [Gkioulekas et al. 2013; Pfeiffer and Marroquim 2012], especially for a dipole rendering system.

To solve this, we approximate R_s using a bilinear interpolation method as in Fig. 11. Specifically, we pre-render and store images with different σ_a and σ_s on a regular 2D grid in log space, covering the potential range of the 2 parameters. Then, for a given output parameter set σ_a , σ_s and w , we first use σ_a and σ_s to perform a bilinear interpolation on the 2D grid to approximate the rendered image and then use w to scale the overall brightness of the image. This way, R_s becomes differentiable and can be evaluated efficiently through a bilinear interpolation. For the training scene, where the model is bounded in a unit radius sphere, we find that a σ_a and a σ_s both in the range $[2^2, 2^8]$ are sufficient. Therefore, we define our post-processing mapping in Equation 12 such that the two parameters both fall in the range. Similarly for w , a range in $[0, 10]$ is sufficient.

At this point, we have all the ingredients to train the neural network. We generated a dataset with random input parameters and corresponding rendered images. In each training iteration, we randomly select a subset of the training examples from the dataset and minimize the loss function using gradient descent. We will provide more details about our training settings in Section 6.

6 IMPLEMENTATION

In this section, we provide key implementation details of two relevant aspects: training and rendering.

Table 2. Sampling strategies used for our data set generation, where u is a uniform random number in range $[0, 1]$. Please refer to Table 1 for the meaning of each parameter. Note that we choose these strategies to cover the range of realistic hair/fur parameters for the hair/fur model in Yan et al. [2017], which we list in the parenthesis next to each parameter.

$\kappa(0.3, 0.9)$	$\sigma_{c,a}(0.01, 4)$	$\sigma_{m,s}(0.25, 4)$	$\eta(1.2, 1.7)$	$l(0.3, 2.5)$
$0.6u + 0.3$	2^{9u-7}	2^{4u-2}	$0.5u + 1.2$	$2.2u + 0.3$

6.1 Neural Network Settings

We generate both the training samples and the 2D grid of output rendered images in Fig. 11 using the *Fur pelt* scene. For training samples, we rendered 1000 images in 128×128 resolution with 1024 samples per pixel using different input parameter sets randomly sampled from the 5D input space. Note that for each image, we use different $\sigma_{c,a}$ in the RGB channels to increase diversity. Figure 12 shows some of the examples. We list our sampling strategies for each input dimension in Table 2. For fixed parameters, we set β_m and β_n to 0.1 and the rest to 0. As introduced in Section 5, these fixed parameters do not have observable impact on scattered components. For the 2D grid of rendered images, we generate a 15×15 uniform grid in the log space of σ_a and σ_s .

We trained our neural network using Tensorflow with Adam gradient descent optimizer. Training takes about 20 minutes to converge on our data set. We obtained 0.90 structural similarity and 0.070 L_1 norm on average in our training set. We also tested the neural network in a validation set, where we generated 100 images using the same sampling strategy. In the validation set, we obtained 0.89 structural similarity index and 0.072 L_1 norm. Note that a structural similarity index around 0.9 suggests fairly high resemblance between two images. We show more validation of our neural network in Section 7.

6.2 Rendering

Parameter Scaling. The dipole model parameters we obtained in Section 5 is for the scale of our training scene. To generalize to scenes with different scales, we use the estimated mean free path to scale σ_a and σ_s , since these 2 parameters are inversely proportional to the mean free path. Specifically, for a given input scene, we estimate the average surface density d of the hair/fur fibers. The mean free path of the hair/fur volume is inversely proportional to \sqrt{d} . This is because if we scale the scene up by a factor of 2, the average distance between 2 closest hair/fur fibers also scales up by 2 and the surface density of the hair/fur fiber becomes 1/4 of the original. Suppose the training scene's estimated hair/fur fiber density is d_0 , we scale both σ_a and σ_s by $\sqrt{d_0/d}$ for rendering the scene. As we show in Section 7, this simple scaling scheme generalizes very well with different scene scales.

Heterogeneous hair/fur parameters. To support heterogeneous hair/fur parameters, we simply perform the parameter conversion using the neural network introduced in Section 5.2 at each shading point and evaluate the dipole model using local parameters. Note that the size of our neural network structure is very small, so that the overhead of the evaluation is minimal compared to the ray tracing cost. We validate this using the hair with medulla scene shown in Fig. 6(e). The hair scene is homogeneous, so the neural network only needs to be evaluated once at the beginning. For comparison,

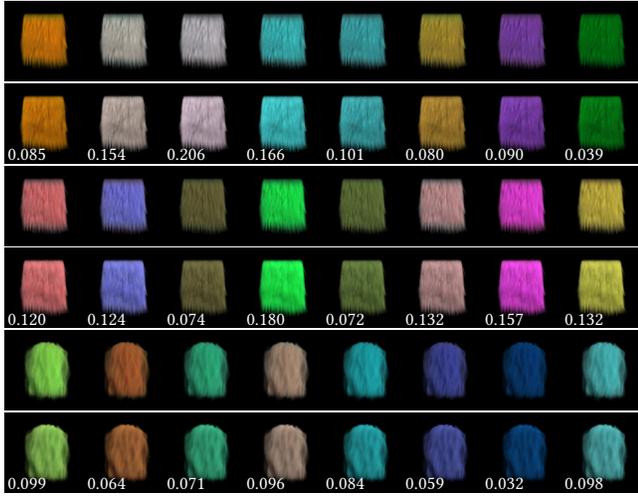


Fig. 12. Validation of our neural network training. Odd numbered rows are rendered with path tracing of actual hair/fur fibers as reference. Even numbered rows are rendered using best fitted/predicted parameters as BSSRDF. Top two rows: selected training set and fitting. Middle two rows: selected test set with same geometry but different lighting conditions. Bottom two rows: selected test set on another scene. Our trained neural network predicts a perceptually similar match, compared to the reference. The numbers in the figures show the loss value of each fitted image. Recall that the loss function is defined in Section 5.3.

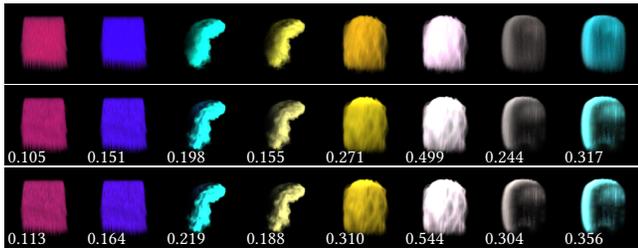
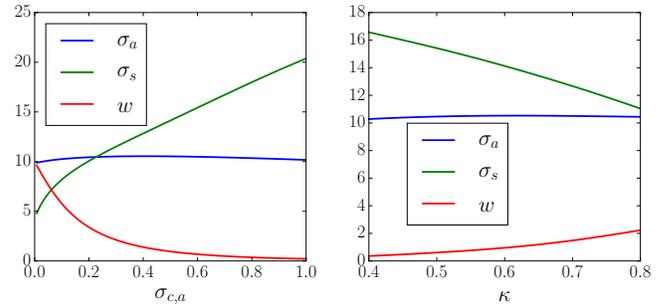


Fig. 13. Comparison of our neural network training using single scene (fur pelt, leftmost two columns) and multiple scenes (all 4 scenes in this figure). The first row is the path traced reference. The second row is predicted using our network trained with a single scene. The third row is predicted using our network trained with multiple scenes. Our neural network predicts a perceptually similar match using single and multiple scenes. The last two examples also show inaccurate cases of our neural network, especially with back lighting in the rightmost two columns.

we render the scene with/without re-evaluating the neural network at each shading point. The render time without re-evaluation is 99s and the render time with re-evaluation is 105s. This suggests that the overhead of the neural network evaluation is less than 10% of the total render time.

Multi-bounce Illumination For multi-bounce illumination with dual scattering, we designed an importance sampling scheme for the dual scattering shading which we provide in our supplemental materials. With the importance sampling scheme, we can follow the traditional path tracing routine by recursively sampling rays and computing shading with dual scattering models. However, one issue is that strictly following the original path tracing routine would overestimate energy. This is because the dual scattering model



(a) Curves showing how the fitted dipole parameters vary with the absorption coefficient in the cortex. (b) Curves showing how the fitted dipole parameters vary with the medulla index.

Fig. 14. Parameter conversion curves learned by our MLPNN.

Table 3. Statistics for all our scenes, rendered in 720p resolution. For each scene, there are # Strands hair/fur fibers, each with # Segs line segments. Each scene is rendered using # Samples using our 3-component global illumination model. Pre-distributing dipole samples usually takes less than 1/10 of the overall rendering time, so we do not list these timings here.

	Fig.	#Strands	#Segs	#Samples	Time
Pelt	1	12.5K	4	42	7.0min
Raccoon	15	260K	22	64	5.4min
Wolf	16	1.9M	5	17	5.8min
Hamster	18	580K	15	31	5.0min
curly	17	53K	64	42	5.0min

already handles local multiple scattering in a neighborhood. This way, if the next bounce is a nearby hair fiber, the energy would be over estimated. To solve this we set a threshold d_m for the distance between 2 consecutive bounces with dual scattering material. If the distance is smaller than d_m , we stop the recursive path tracing since the next bounce has already been considered by the dual scattering shading at the current bounce, otherwise we continue tracing the ray as usual. We apply the same strategy for the recursive irradiance sampling in the dipole model for the same reason. For all our test scenes, we set $d_m = 100r$, where r is the average hair/fur fiber radius of the hair/fur model at the current bounce point.

7 RESULTS

In this section, we first validate our trained neural network, then show rendering results with full global illumination generated using our BSSRDF model, and compare them with previous work. We implement our model in the Mitsuba renderer [Jakob 2010], and generate all results using an Intel 6-core i7 4960X CPU, hyperthreaded to 12 threads. The source code for both neural network training and BSSRDF model implementation, as well as our trained neural network are available on <http://viscomp.ucsd.edu/projects/furbssrdf>.

Validation of the training. As a typical shape of fur growth, the *Fur pelt* scene is used to train the neural network. We first validate the trained neural network with the same scene geometry but different lighting conditions and different fur fiber parameters. Then we validate using a new hair geometry, also with different lighting and fiber parameters, to demonstrate that it is sufficient to train on



Local+Unscat. Scattered Full model Path (ES) Path (REF)
64spp, 2.5min 64spp, 4.1min 64spp, 5.1min 64spp, 2.5min 1Kspp, 40min

Fig. 15. (Top row) A rendering of the *Raccoon* model using our efficient 3-component global illumination model. (Bottom row) Insets showing decomposed renderings of different components, compared to path traced reference with 64 samples and 1024 samples per pixel, respectively. We find that the main source of noise is the unscattered component using dual scattering. The scattered component converges much faster. Overall, our full model produces much less noise than path tracing for equal samples.

only one scene and use the result to predict other scenes. Figure 12 shows some of the validation results. We also include all validations in the supplemental material. Figure 13 demonstrates that using one scene to train our neural network is enough. Using multiple scenes only results in slightly better loss, and it doesn't help with difficult cases either. Slight color differences with the reference still exist. And back lighting is still inaccurate to represent using the simple dipole model. In the future, more accurate models such as extended dipole [Donner and Jensen 2008] and multi-pole [Donner and Jensen 2005] may help with these cases.

Figure 14 shows curves learned by the MLPNN structure. In Fig. 14 (a), we show how the dipole parameters vary with the absorption coefficient in the cortex, while keeping other input parameters fixed. As $\sigma_{c,a}$ increases, each hair fiber absorbs more energy. This would result in a darker and sharper looking appearance of the hair model. In this case, we can see that the dipole model decreases w to match the darker appearance and increases σ_s to match the sharper look. In Fig. 14 (b), we show how the dipole parameters vary with the medulla index, while keeping other input parameters fixed. Increasing the medulla size would produce more scattering events between hair fibers, which leads to a brighter and more smooth looking appearance. The learned dipole model responds by increasing w and decreasing σ_s in this case, which produces similar effects. Note that σ_a stays relatively flat in the learned function. This is probably because some correlations exist among the 3 output parameters. For example, increasing σ_a and decreasing w would both decrease the brightness of the rendered model. For this reason, the MLPNN probably finds it easier to fix one of them in the learning process. However, we should note that the learned relatively constant values

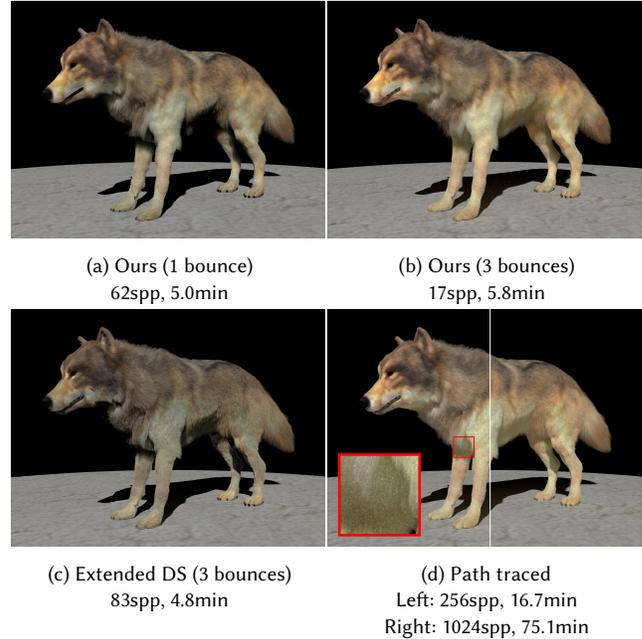


Fig. 16. The *Wolf* scene rendered using a point light. We compare our method with 1 bounce and 3 bounces with the extended dual scattering method with 3 bounces and the path traced reference. Our method with 1 bounce is already better than the extended dual scattering method in terms of more accurate color and softer appearance. Moreover, our method with 3 bounces enables inter-reflections between fibers in the same fur volume, filling shadows on the belly and the limbs and producing a similar appearance as the reference, but still achieves a minimum of 3× speed up. The path tracing with 256 spp is still noisy, as the inset shows.



Fig. 17. The *Curly hair* scene rendered with a point light. We compare our method (b) with the extended dual scattering method (a) and Yan et al. [2017] (c). Although the overall color/intensity is generated by the extended dual scattering, it completely fails to capture the correct color distribution. The highlight is shifted towards the bottom left. In contrast, our method is able to produce a much more accurate appearance, as compared with the ground truth, yet is still an order of magnitude faster.

of σ_a are still crucial to the appearance of the dipole model and all 3 parameters are essential.

Scene configurations. After validation of our neural network training, we now use the trained neural network to convert parameters for actual scenes. Most of our scenes are taken from Yan et al. [2017] to enable direct comparison to their method. We list all the scene configurations including geometry complexity and performance



Fig. 18. The *Hamster* scene rendered using a spot light. The extended dual scattering method with 3 bounces still has the color mismatching issue. However, our method with 3 bounces closely resembles the path traced reference, but requires many fewer samples and is an order of magnitude faster. Significant noise can be observed for equal time path tracing when zooming in. The higher-order bounces result in an overhead for our method, since further bounces also require significant dipole queries. However, since the dipoles' contributions are smooth, even with the reduced sampling rate, our method still converges fast. The photon mapping approach produces more noise than ours in equal time, and converges even slower than path tracing.

in Table 3. We compare the scenes with our extended dual scattering and Yan et al. [2017] with path-traced global illumination as reference, and we describe them next. In all our scenes, we use 3 recursive bounces for our dipole approximation and 4 indirect samples for irradiance estimation.

Fur pelt. We render the *Fur pelt* with a blocking sphere casting a shadow onto it, and compare our rendering with dual scattering. As demonstrated in Fig. 1, dual scattering only lights the un-blocked regions, resulting in hard shadow boundaries and no color bleeding. However, significant scattering effects are observed in the reference, and our BSSRDF model matches that much more accurately.

Raccoon. The *Raccoon* scene is rendered with a uniform sky environment map with a sharp point approximating the sun. Figure 15 shows decomposed components and corresponding rendering time. We can see that, most energy is captured by the BSSRDF part, which soon becomes noiseless as the number of samples increases, thanks to the smoothness of the dipoles' contribution.

Wolf. The *Wolf* scene is to show our dipole model with multiple global illumination bounces. In Fig. 16, we show side by side comparisons with and without multiple bounces using our method, rendered using a point light, and compare with the extended dual scattering, roughly for the same rendering time (~ 5 min). The extended dual scattering considers multiple bounces, but still leaves hard and dark shadows around regions near the belly and the limbs. Furthermore, it has a significant color difference with the reference, and it generates hard and solid appearance, especially on the head. Our method with 1 bounce can already capture accurate color and soft appearance. With 3 bounces, we are able to generate similar appearance with the reference. The rendering time increases roughly linearly with the number of bounces, but we are still at least 3 times faster, since the path traced result at 256 spp is still noisy when zooming in.

Hamster. The *Hamster* scene is rendered using a spot light in Fig. 18 with increasing numbers of bounces. We also compare with the extended dual scattering method with 3 bounces, and we can see that the dense fur fibers soon diminish global scattering from dual

scattering, resulting in an overly-dark appearance. Conversely, our method with 3 bounces matches the reference much closer and is still fast. The photon mapping approach [Moon and Marschner 2006] uses more photons (100M) than the paths in the path traced reference for equal quality comparison with our method. And their 6D photon query becomes inefficient for fur, since the fur fibers spread much wider than hair in the angular space. So, the performance of the photon mapping approach is significantly worse than ours.

Furthermore, in the accompanying video, we rotate the camera and demonstrate that our model doesn't have to re-generate all dipole samples over frames, as long as the relative lighting condition of the model doesn't change. We also include a moving light video with fixed and dynamically sampled dipoles over frames. With fixed samples, there is almost no flickering. Moreover, even with dynamically sampled dipoles, the flickering is arguably tolerable. For animated/deformable objects, we believe correlated sample locations would help, and we leave it to future work.

Curly hair. Our trained neural network can also predict the appearance of hair. We compare renderings of *Curly hair* using our model with medullary index $\kappa = 0.36$ against dual scattering and the path traced reference, as shown in Fig. 17. It demonstrates that dual scattering fails to capture the actual highlight positions from scattering, and it always produces a flat shading distribution.

Note that, since most of our results contain spatially-varying colors and various fur orientations and densities, our method is capable of handling heterogeneity as a general BSSRDF solution. We believe that specific complex BSSRDF models, such as multipole [Donner and Jensen 2005] and quantized diffusion [D'Eon and Irving 2011] will further improve it, and we leave them as future work.

8 CONCLUSION AND FUTURE WORK

We have presented the first scattering model to efficiently approximate global illumination within hair and fur volumes. We analyze and point out failure cases of classic dual scattering, and organize our model as three components: direct illumination, dual scattering for unscattered lobes, and BSSRDF for all other scattering events.

We convert properties from hair and fur fibers to BSSRDF parameters by training a neural network on only one scene. Our model supports various lighting conditions including environment maps, and enables hair to hair inter-reflections. We show close matches of our predicted renderings using the dipole method, compared with the path traced reference of Yan et al. [2017].

In the future, one natural extension to our method is to make it real-time, since both the dual scattering and the dipole have been demonstrated as suitable for efficient GPU implementation [d'Eon et al. 2007]. Another practical direction is to support SRBF lights [Ren et al. 2010; Xu et al. 2011] to enable real-time hair and fur rendering under environment illumination. Furthermore, it is also worth exploring ways to describe the BSSRDF more accurately, such as learning the diffusion profiles directly, or using anisotropic dipoles to account for hair and fur fibers' strong local similarity in their orientations.

ACKNOWLEDGMENTS

We would like to thank the primary reviewer for their detailed and insightful comments and discussions with us, which led to a substantial improvement of this work. This research was funded in part by NSF grants 1451828 and 1703957, an NVIDIA fellowship and the UC San Diego Center for Visual Computing.

REFERENCES

- James F. Blinn. 1982. Light Reflection Functions for Simulation of Clouds and Dusty Surfaces. *SIGGRAPH Comput. Graph.* 16, 3 (July 1982), 21–29. DOI: <https://doi.org/10.1145/965145.801255>
- Matt Jen-Yuan Chiang, Benedikt Bitterli, Chuck Tappan, and Brent Burley. 2016. A Practical and Controllable Hair and Fur Model for Production Path Tracing. *Computer Graphics Forum* 35, 2 (2016), 275–283. DOI: <https://doi.org/10.1111/cgf.12830>
- Eugene d'Eon, Guillaume Francois, Martin Hill, Joe Letteri, and Jean-Marie Aubry. 2011. An Energy-conserving Hair Reflectance Model. In *ACM Transactions on Graphics (TOG)*. 1181–1187. DOI: <https://doi.org/10.1111/j.1467-8659.2011.01976.x>
- Eugene D'Eon and Geoffrey Irving. 2011. A Quantized-diffusion Model for Rendering Translucent Materials. In *ACM SIGGRAPH 2011 Papers (SIGGRAPH '11)*. ACM, New York, NY, USA, Article 56, 14 pages. DOI: <https://doi.org/10.1145/1964921.1964951>
- Eugene d'Eon, David Luebke, and Eric Enderton. 2007. Efficient Rendering of Human Skin. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques (EGSR'07)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 147–157. DOI: <https://doi.org/10.2312/EGWR/EGSR07/147-157>
- Eugene d'Eon, Steve Marschner, and Johannes Hanika. 2013. Importance Sampling for Physically-based Hair Fiber Models. In *SIGGRAPH Asia 2013 Technical Briefs*. Article 25, 4 pages. DOI: <https://doi.org/10.1145/2542355.2542386>
- Craig Donner and Henrik Wann Jensen. 2005. Light Diffusion in Multi-layered Translucent Materials. *ACM Trans. Graph.* 24, 3 (July 2005), 1032–1039. DOI: <https://doi.org/10.1145/1073204.1073308>
- Craig Donner and Henrik Wann Jensen. 2008. Rendering Translucent Materials Using Photon Diffusion. In *ACM SIGGRAPH 2008 Classes (SIGGRAPH '08)*. ACM, New York, NY, USA, Article 4, 9 pages. DOI: <https://doi.org/10.1145/1401132.1401138>
- Julie Dorsey, Alan Edelman, Henrik Wann Jensen, Justin Legakis, and Hans Köhling Pedersen. 1999. Modeling and Rendering of Weathered Stone. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '99)*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 225–234. DOI: <https://doi.org/10.1145/311535.311560>
- Eric Enderton, Erik Sintorn, Peter Shirley, and David Luebke. 2011. Stochastic transparency. *IEEE transactions on visualization and computer graphics* 17, 8 (2011), 1036–1047.
- Jeppé Revall Frisvad, Toshiya Hachisuka, and Thomas Kim Kjeldsen. 2014. Directional Dipole Model for Subsurface Scattering. *ACM Trans. Graph.* 34, 1, Article 5 (Dec. 2014), 12 pages. DOI: <https://doi.org/10.1145/2682629>
- Ioannis Gkioulekas, Shuang Zhao, Kavita Bala, Todd Zickler, and Anat Levin. 2013. Inverse volume rendering with material dictionaries. *ACM Transactions on Graphics (TOG)* 32, 6 (2013), 162.
- Eric Heitz, Jonathan Dupuy, Cyril Crassin, and Carsten Dachsbacher. 2015. The SGGX Microflake Distribution. *ACM Trans. Graph.* 34, 4, Article 48 (July 2015), 11 pages. DOI: <https://doi.org/10.1145/2766988>
- Christophe Hery and Ravi Ramamoorthi. 2012. Importance sampling of reflection from hair fibers. *Journal of Computer Graphics Techniques (JCGT)* 1, 1 (2012), 1–17.
- Wenzel Jakob. 2010. Mitsuba renderer. <http://www.mitsuba-renderer.org>. (2010).
- Wenzel Jakob, Adam Arbree, Jonathan T. Moon, Kavita Bala, and Steve Marschner. 2010. A Radiative Transfer Framework for Rendering Materials with Anisotropic Structure. *ACM Trans. Graph.* 29, 4, Article 53 (July 2010), 13 pages.
- Henrik Wann Jensen and Juan Buhler. 2002. A rapid hierarchical rendering technique for translucent materials. In *ACM Transactions on Graphics (TOG)*, Vol. 21. ACM, 576–581.
- Henrik Wann Jensen, Justin Legakis, and Julie Dorsey. 1999. Rendering of Wet Materials. In *Proceedings of the 10th Eurographics Conference on Rendering (EGWR '99)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 273–282. DOI: <https://doi.org/10.2312/EGWR/EGWR99/273-282>
- Henrik Wann Jensen, Stephen R Marschner, Marc Levoy, and Pat Hanrahan. 2001. A practical model for subsurface light transport. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. ACM, 511–518.
- Pramook Khungurn, Daniel Schroeder, Shuang Zhao, Kavita Bala, and Steve Marschner. 2015. Matching Real Fabrics with Micro-Appearance Models. *ACM Trans. Graph.* 35, 1, Article 1 (Dec. 2015), 26 pages. DOI: <https://doi.org/10.1145/2818648>
- Tom Lokovic and Eric Veach. 2000. Deep Shadow Maps. In *SIGGRAPH 00*. 385–392.
- Stephen R. Marschner, Henrik Wann Jensen, Mike Cammarano, Steve Worley, and Pat Hanrahan. 2003. Light Scattering from Human Hair Fibers. *ACM Transactions on Graphics (TOG)* 22, 3 (2003), 780–791. DOI: <https://doi.org/10.1145/882262.882345>
- Johannes Meng, Marios Papas, Ralf Habel, Carsten Dachsbacher, Steve Marschner, Markus Gross, and Wojciech Jarosz. 2015. Multi-Scale Modeling and Rendering of Granular Materials. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 34, 4 (July 2015). DOI: <https://doi.org/10.1145/2766949>
- Jonathan T. Moon and Stephen R. Marschner. 2006. Simulating Multiple Scattering in Hair Using a Photon Mapping Approach. *ACM Trans. Graph.* 25, 3 (2006), 1067–1074. DOI: <https://doi.org/10.1145/1141911.1141995>
- Thomas Müller, Marios Papas, Markus Gross, Wojciech Jarosz, and Jan Novák. 2016. Efficient Rendering of Heterogeneous Polydisperse Granular Media. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 35, 6 (December 2016), 168:1A–168:14. DOI: <https://doi.org/10.1145/2980179.2982429>
- Gustavo Pfeiffer and Ricardo Marroquim. 2012. BRDF Fitting Using Inverse Global Illumination and Stochastic Optimization. (2012).
- Matt Pharr and Pat Hanrahan. 2000. Monte Carlo Evaluation of Non-linear Scattering Equations for Subsurface Reflection. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '00)*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 75–84. DOI: <https://doi.org/10.1145/344779.344824>
- Zhong Ren, Kun Zhou, Tengfei Li, Wei Hua, and Baining Guo. 2010. Interactive Hair Rendering Under Environment Lighting. In *ACM SIGGRAPH 2010 Papers (SIGGRAPH '10)*. ACM, New York, NY, USA, Article 55, 8 pages. DOI: <https://doi.org/10.1145/1833349.1778792>
- Iman Sadeghi and Rasmus Tamstorf. 2010. *Efficient Implementation of the Dual Scattering Model in RenderMan*. Technical Report. Tech. rep., Walt Disney Animation Studios.
- Erik Sintorn and Ulf Assarsson. 2008. Real-time approximate sorting for self shadowing and transparency in hair rendering. In *Proceedings of the 2008 symposium on Interactive 3D graphics and games*. ACM, 157–162.
- Erik Sintorn and Ulf Assarsson. 2009. Hair Self Shadowing and Transparency Depth Ordering Using Occupancy Maps. In *Symposium on Interactive 3D Graphics and Games*. 67–74. DOI: <https://doi.org/10.1145/1507149.1507160>
- Zhou Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* 13, 4 (April 2004), 600–612. DOI: <https://doi.org/10.1109/TIP.2003.819861>
- Kun Xu, Li-Qian Ma, Bo Ren, Rui Wang, and Shi-Min Hu. 2011. Interactive Hair Rendering and Appearance Editing under Environment Lighting. *ACM Transactions on Graphics* 30, 6, Article 173 (2011), 173:1–173:10 pages.
- Lingqi Yan, Chiwei Tseng, Henrik Wann Jensen, and Ravi Ramamoorthi. 2015. Physically-accurate fur reflectance: modeling, measurement and rendering. *ACM Transactions on Graphics (TOG)* 34, 6 (2015), 185.
- Ling-Qi Yan, Henrik Wann Jensen, and Ravi Ramamoorthi. 2017. An Efficient and Practical Near and Far Field Fur Reflectance Model. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2017)* 36, 4 (2017).
- Xuan Yu, Jason C. Yang, Justin Hensley, Takahiro Harada, and Jingyi Yu. 2012. A Framework for Rendering Complex Scattering Effects on Hair. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (I3D '12)*. ACM, New York, NY, USA, 111–118. DOI: <https://doi.org/10.1145/2159616.2159635>
- Cem Yuksel and John Keyser. 2008. Deep Opacity Maps. *Computer Graphics Forum* 27, 2 (2008), 675–680. DOI: <https://doi.org/10.1111/j.1467-8659.2008.01165.x>
- Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2011. Building Volumetric Appearance Models of Fabric Using Micro CT Imaging. *ACM Trans. Graph.* 30, 4, Article 44 (July 2011), 10 pages. DOI: <https://doi.org/10.1145/2010324.1964939>
- Shaung Zhao, Lifan Wu, Frédo Durand, and Ravi Ramamoorthi. 2016. Downsampling Scattering Parameters for Rendering Anisotropic Media. *ACM Trans. Graph.* 35, 6 (2016).
- Arno Zinke, Cem Yuksel, Andreas Weber, and John Keyser. 2008. Dual Scattering Approximation for Fast Multiple Scattering in Hair. *ACM Transactions on Graphics* 27, 3, Article 32 (2008), 10 pages. DOI: <https://doi.org/10.1145/1360612.1360631>